

USER REQUEST EVALUATION TOOL (URET) CONFLICT PROBE PERFORMANCE AND BENEFITS ASSESSMENT

D. J. Brudnicki and A. L. McFarland

The MITRE Corporation

Center for Advanced Aviation System Development

7525 Colshire Drive, M.S. W273

McLean, VA USA 22102-3481

© 1997 The MITRE Corporation.

This is the copyright work of The MITRE Corporation, and was produced for the U. S. Government under Contract Number DTFA01-93-C-00001, and is subject to Federal Acquisition Regulation Clause 52.227-14, Rights in Data—General, Alt. III (JUN 1987) and Alt. IV (JUN 1987). No other use other than that granted to the U. S. Government, or to those acting on behalf of the U. S. Government, under that Clause is authorized without the express written permission of The MITRE Corporation. For further information, please contact The MITRE Corporation, Contracts Office, 1820 Dolley Madison Blvd., McLean, VA 22102, (703) 883-6000.

The contents of this material reflect the views of the authors. Neither the Federal Aviation Administration nor the Department of Transportation makes any warranty or guarantee, or promise, expressed or implied, concerning the content or accuracy of the views expressed herein.

ABSTRACT

In order to respond to rising demands for services from the National Airspace System (NAS), the Federal Aviation Administration (FAA) and The MITRE Corporation's Center for Advanced Aviation System Development (CAASD) have developed and evaluated a set of en route air traffic control (ATC) decision support capabilities based on many of years of earlier work on the Advanced En Route ATC (AERA) program. The component capabilities embodied in a prototype referred to as the User Request Evaluation Tool (URET) include a conflict probe that continuously checks current flight plan trajectories for strategic conflicts, and a Trial Planning function that allows a controller to expeditiously evaluate problem resolutions before they are issued as clearances. These capabilities are intended to provide the flexibility and decision support needed to allow more user preferences to be met while continuing to maintain or enhance today's level of safety.

The URET prototype was deployed to the Indianapolis Air Route Traffic Control Center (ARTCC) for field trials in January 1996. This paper presents key findings from the field trials and related activities that have occurred since. It includes a summary of the evaluation results, a detailed algorithmic performance assessment of URET's conflict probe and component capabilities, and a discussion of the projected benefits that can be expected with a national deployment of the system. The field trials have strongly indicated thus far that URET capabilities are operationally acceptable for both planning and clearance decision making purposes at the en route sector. The algorithmic analyses quantify those results in terms of conflict likelihood and warning time to effectively define the performance levels needed to support that degree of utility and acceptability. Finally, the benefits analyses focus on contrasting today's operations with expected increases in system flexibility, efficiency, and utilization in terms of projected annual savings to the aviation community.

acknowledgments

The authors wish to acknowledge the contributions to this paper of several of their CAASD colleagues. Bill Arthur conducted the Conflict Likelihood Statistical Analysis and Ken Lindsay conducted the Conflict Likelihood Voice Tape Analysis. Mike White and Bill Kuhn conducted the nationwide quantitative NRP and Conflict Probe Benefits study. Susan Schultheis led the field evaluation activity and wrote results of that activity from which some of this material has been extracted. We also wish to recognize the valuable cooperation, support, and suggestions we have received from the controllers and managers at Indianapolis ARTCC throughout the development and field evaluation of URET.

Introduction

Background

Current ATC operations in many areas of the United States are highly structured and restrictive. Given the tools that are currently available to en route controllers, the structure and restrictive nature of operations is necessary in order to manage air traffic within an acceptable level of safety. It is, however, in conflict with users' desires for more flexibility, i.e., the freedom to fly more preferred routes and altitudes from origin to destination. The airspace user community has expressed concerns that the current system is unable to keep pace with this growing demand. Acknowledging the validity of these concerns, the FAA asked CAASD to develop and demonstrate in field use tools that could be implemented in the near term as an adjunct to the current operational computer system.

There is consensus among the service provider and service user communities that, at a minimum, controllers require earlier notice of developing traffic problems to allow them to manage current and projected workloads more efficiently. In order to meet this need, the FAA and CAASD have been conducting extensive research on decision support tools that provide for the timely detection and resolution of predicted problems. The URET prototype currently being evaluated at the Indianapolis ARTCC provides those capabilities. The specific objectives of the Indianapolis field evaluation are:

- 1) Validate historical lab-tested concepts within a real-world context. These concepts were developed as part of the AERA program over a number of years, in close collaboration with field controllers. URET provides a subset of those AERA capabilities.
- 2) Define the initial set of acceptable and beneficial decision support capabilities for implementation. The evaluation not only helps determine the potential ATC and user benefits that can be accrued, but reduces the risk associated with the acquisition and deployment of those capabilities on a larger scale.

The URET Prototype

URET processes real-time flight plan and track data from the ARTCC Host computer through a one-way interface. These data are combined with site adaptation, aircraft performance characteristics, and winds and temperatures from the National Weather Service in order to build four-dimensional flight profiles, or trajectories, for all flights within or inbound to the facility. URET also provides a reformance function that adapts each trajectory to the observed speed, climb rate, and descent rate of the modeled flight. For each flight, incoming track data are continually monitored and compared to the trajectory in order to keep it within acceptable tolerances.

URET maintains current plan trajectories, i.e., those that represent the current set of flight plans in the system, and uses them to continuously check for conflicts. When a conflict is detected, URET determines which sector to notify and displays an alert to that sector up to 20 minutes prior to the start of that conflict. URET also provides a trial plan function. Trial planning allows a controller to check a desired flight plan amendment for potential conflicts before a clearance is issued.

These capabilities are packaged within a Computer Human Interface (CHI) that includes text and graphic information. The text-based Aircraft List and Plans Display manage the presentation of current plans, trial plans, and conflict probe results for each sector. The Graphic Plan Display provides a graphical capability to view aircraft routes and altitudes, predicted conflicts, and trial plan results. In addition, the point-and-click interface enables quick entry and evaluation of trial plan route, altitude, or speed changes. Finally, the Wind Grid display provides a visual representation of forecast winds and temperatures at selected altitudes.

FIELD TRIALS

Evaluation Overview

As of April 1997, a total of 70 hours of evaluation have been completed on the control room floor of the Indianapolis ARTCC. A total of 15 Full Performance Level (FPL) controllers and 5 supervisors participated at several high, super-high, and low altitude sectors under varying traffic conditions. Two Traffic Management Specialists participated in training but were not directly involved in the evaluation sessions.

Prior to formal evaluations, each participant received a total of 16 hours of training. This was comprised of classroom instruction, hands-on training using scripted scenarios, and practice using recorded or live scenarios. Training played a key role in the overall process. The classroom instruction provided the details about system functionality, operational concepts and procedures, and the role of the participants in the evaluations themselves; hands-on operation promoted the necessary

familiarization with the user interface; and training in general facilitated initial acceptance of the tool.

Evaluations commenced in February 1996. Initially, a team of 3 trained controllers worked at a sector with URET during scheduled evaluation sessions. The team consisted of the R-side controller, D-side controller, and URET operator. The URET workstation was positioned on a cart adjacent to the sector such that the controller could use URET and observe sector operations simultaneously. In this configuration, the R and D controllers were, as an operational rule, unable to interact with the URET operator. But by April, this rule was dropped, and the operator was able to communicate URET information to the sector team to be considered in the decision-making process. By December, the facility was able to integrate the URET workstation into the D side at one sector for evaluation sessions. In this configuration, the sector team no longer included a separate URET operator. The D controller subsumed this role. Finally, by April of 1997, a trial of continuous use operations was conducted. The prototype was operated in an integrated D-side configuration, that included limited flight strips, for a complete shift (8 hours) on three consecutive days.

Evolution of Capabilities

In 1996, three major URET software deliveries were made to the Indianapolis ARTCC. In January, Delivery 1 (or D1) included a passive probe function, whereby conflicts were displayed only upon controller request. In May, the next delivery (D1.1) provided a continuous probe capability for current plans and continuous probing of selected trial plans (an Automated Replan function). D1.1 did not, however, include sector notification logic. All sectors were shown all alerts immediately upon detection. In November, version D2.0 included logic to notify the sector only of alerts that are predicted to occur in that sector up to 20 minutes into the future. Furthermore, logic was incorporated to compute and use conflict probability to determine notification time. The intent is to delay notification of low probability conflicts since they are more likely to be reconfirmed and, thus, re-evaluated in the time before notification.

Current plans and activities include D2.1 and D3.0 versions of URET. The major additional capabilities in D2.1 are automated coordination of trial plans between sectors, and CHI enhancements. It also includes a make-current function that allows a trial plan to be sent to the Host as a flight plan amendment. This, however, is not currently supported by the Host interface and is expected to be evaluated some time beyond the delivery of the D3.0 system. D3.0 focuses on interfacility capabilities. It enables two or more URET systems in adjacent facilities to communicate information in order to provide more accurate and timely predictions, and to allow for more seamless and strategic operations among those facilities.

Results and Analysis

A thorough debrief and detailed analysis was conducted at the conclusion of each evaluation session. The participants were asked a number of questions concerning the accuracy of the probe, the acceptability and utility of the capabilities, the potential for benefits, etc. What follows are the results of those discussions.

To date, controllers are unanimous in their support of the tool. With respect to the quality of information presented, the trajectory modeling and conflict prediction functions are considered to be operationally accurate and reliable, and are believed to be suitable for both clearance decision making and sector planning. The presentation of information and the method of interacting with the system (the CHI) was found to be efficient and operationally acceptable.

With respect to benefits, the capabilities are expected to be appropriate for use at all sectors, both high and low. They are anticipated to yield the following benefits to the sector team:

- 1) Enhanced safety from early, accurate warnings.
- 2) Enhanced productivity from the D controller (with a potential for less workload given reduced requirements for flight strip manipulation and elimination of the need to scan strips for potential conflicts) with respect to planning, strategic problem detection, and problem resolution.
- 3) Improved quality of flight plan and strategic problem information to the sector as a whole.
- 4) Potential for reduced coordination between sectors when plans are known to be problem-free.
- 5) Better team relationship between the D- and R-side.

- 6) Capabilities that do not degrade as traffic complexity and volume increase.

With respect to airspace users, the following benefits are anticipated:

- 1) Overall enhanced margin of safety.
- 2) Separation based on boundary crossing restrictions could be relaxed given that controllers would have problem detection and separation assistance. Thus, aircraft could fly at more preferred altitudes.
- 3) Potential for longer direct or wind-optimal routings.
- 4) More efficient maneuvers, with fewer secondary maneuvers required since the trial planning function allows clearances to be checked for conflicts before they are issued.

The Benefits Analysis section of this paper provides a more in-depth evaluation of benefits from some additional studies that were done to supplement these field trial results.

FUNCTIONAL PERFORMANCE

Approach

The functional performance of the URET conflict probe is quantified using two distinct methods. The first method is a combined empirical and analytical evaluation of the URET trajectories and predicted conflicts using a real-world traffic scenario. Scenario data are obtained from the ARTCC in the form of System Analysis Recordings (SAR). Flight plan and track report messages are extracted from the SAR and processed by the URET trajectory modeling, track management, and conflict probe functions just as they would be in real-time. During this processing, key data (trajectory segments, predicted conflicts, etc.) are collected. When the scenario run is complete, a combination of commercial and CAASD-developed tools are used to visually analyze specific flight profiles, as well as calculate the aggregate functional performance results. Key to this approach is the use of the actual field recordings. It assures that the reported results represent the true performance of the system within an operational ATC context.

The second method of quantifying functional performance is based on an analysis of the clearances that were issued during the scenario time period. The clearances are transcribed from the voice tapes for a selected set of sectors and correlated to the conflicts that URET would have predicted had it been in operation at the time. Predicted conflicts are categorized as (1) resulting in a controller resolution action, (2) reversing their conflict status (conflict to no-conflict) prior to a controller resolution action, or (3) resulting in no resolution action. For conflicts in the first category, the distribution of times is calculated for the difference between conflict notification and controller action. This second method of quantifying performance establishes a more comprehensive view of the operational environment and conflict events; however, it is also more time consuming given the requirement to transcribe voice tapes.

Performance Metrics

The evaluation of conflict probe results uses the following two metrics to characterize functional performance:

- 1) *Conflict Likelihood*. When URET notifies the controller of a conflict, what is the probability that it will occur if no action is taken?
- 2) *Conflict Warning Time*. How much warning time does URET provide between conflict notification and the predicted start of the conflict?

Conflict Likelihood addresses the reliability of the probe information while Conflict Warning Time addresses the timeliness of that information. Both affect the degree to which system benefits can be realized. A probe that does not achieve reliable performance may lead to unnecessary maneuvers, or it may be effectively used only in a limited number of situations. With respect to warning time, longer lead times not only support better decision making but enhance system safety.

Results and Analysis

Scenario Characteristics The air traffic scenario used for this analysis is derived from 6 hours of SAR and voice tapes from the Indianapolis ARTCC. The sample was collected during a peak traffic period (1800-0000Z). During this period, the facility handled roughly 1,900 flight entries per hour. With respect to the operational complexity of the sample, it represents a good mix of normal and atypical operations. (Severe Weather (SWAP) routes and Miles-in-Trail (MIT) restrictions were in effect during a portion of the scenario.)

A total of 3,336 flights were extracted from the SAR, of which 2,266 (68%) are commercial transports, 1,003 (30%) are general aviation, and 67 (2%) are military. Of these, roughly 30% were eliminated from the sample for the following reasons:

- 1) 309 flights have track data but no flight plan data. Since a trajectory cannot be built without a flight plan, these flights are not subject to any of the performance metrics. These flights are simply an artifact caused by starting the analysis at the beginning of the traffic sample without the usual 30-minute start-up period used in the field system.
- 2) 708 flights have flight plan data but no track data. The subset of these that are inbound from adjacent facilities are eligible to be probed, and all resulting conflict data are included in the performance measurements. Given the absence of track data, however, the underlying track/trajectory deviation measurements cannot be taken. These are eliminated to remove the artifact of terminating the analysis at the end of the traffic sample.
- 3) 63 flights have too limited a set of track data. The data are limited either in number or by reasonableness of the constituent track report information.
- 4) 15 flights are excluded because they entered a holding pattern without a corresponding Hold message entry.
- 5) 5 flights are excluded from the sample because of incorrectly adapted fixes, forcing a large, unrealistic divergence in the trajectory from the flight path actually flown.

The equipment characteristics for the remaining 2,236 flights are as follows:

- 1) 64% are equipped with area navigation.
- 2) The top ten aircraft types constitute 36% of the sample (greatest to least): DC-9, B-727, B-737-200, MD-80, B-737-300, B-757, Canadair CL-601, B-767, Piper Cherokee Arrow, and Fokker 100.

Finally, the ARTCC Host message composition is as follows:

- 1) 7,180 Flight Plan Entries or Amendments
- 2) 4,349 Interim Altitude Assignments
- 3) 518,485 Track Reports

Conflict Likelihood: Statistical Analysis In the field trials, controllers gave high overall ratings to the validity of displayed conflicts. In addition, they preferred to distinguish conflicts that are more likely to result in a loss of separation from those that are less likely. To distinguish between these cases, conflicts are color-coded. Conflicts with a predicted minimum horizontal separation distance between trajectory centerlines less than or equal to 5 nm are coded in red and those with a predicted minimum separation distance greater than 5 nm are coded in yellow. Thus, one can refer to these as *ired* and *iyellow* conflicts. Controllers considered red conflicts to be of a higher priority than yellow conflicts, although they believed that yellow conflicts nevertheless provide useful information for decision making. The analysis below quantifies probe performance for these two classes of conflicts.

The sector controller team is frequently interested in cases where the aircraft pass close to each other as well as those that

may violate the separation standard. Therefore, the performance analysis below reports conflict likelihood for miss criterion distances of 5, 7, and 9 nm. Also, since the scenario consists of operational data, constituent events include controller actions to maintain separation. Thus, a statistical model is used with the specific conflict data (e.g., encounter geometries, speeds, etc.) to estimate the aggregate likelihood value for a given miss criterion distance. The model is constructed using measured lateral and longitudinal deviation distributions that are dependent on navigational equipment, lookahead time, and number of previous reconformances in each dimension.

For the given scenario, 11,001 total red conflicts and 14,865 total yellow conflicts were predicted. Of those, 23 red conflicts and 1,363 yellow conflicts with a duration of less than 9 seconds were excluded. Figure 1 shows the resultant likelihood values for red and yellow conflicts. It shows, for example, that 78% of the red conflicts predicted by URET would result in a minimum miss distance within 9 nm if no action is taken.




Figure 1. Conflict Likelihood for Red and Yellow Conflicts Using Statistical Model of Lateral and Longitudinal Deviation

The above data describe conflict likelihood for horizontal miss distances. During altitude transitions, vertical errors must also be considered. Given a horizontal miss criterion distance of 7 nm and a vertical miss criterion distance of 1,500/3,000 ft (below/above Flight Level 290), conflict likelihood is 0.45 for red conflicts and 0.24 for yellow conflicts. These results reflect effects such as interim altitudes and uncertainty in top of climb and descent points.

Consistent with controller perceptions, red conflicts are more likely to result in a minimum separation within the miss criterion distance if no action is taken. Most red problems have a high likelihood that a clearance will be needed to maintain separation. And many yellow conflicts will require close monitoring by the sector team. These observations are validated not only by the field trials but by the voice tape analysis described next.

Conflict Likelihood: Voice Tape Analysis Route, altitude, and speed clearances were transcribed from the voice tapes for 3 sectors. Five hours from each sector were transcribed. This served two purposes. First, the clearances were correlated to SAR messages in order to better understand the degree to which these messages reflect the clearances actually issued to the aircraft. Second, the clearances were correlated to conflicts URET would have predicted for this scenario using the approach described earlier.

With respect to the correlation between clearances and Host messages, the following was found:

- 1) A total of 327 route clearances were issued, of which only 59 (18%) were entered as route amendment messages.
- 2) Almost all altitude clearances were correlated to a flight plan amendment or interim altitude message.

3) Of the 33 total speed clearances, none were entered into the Host.

When the scenario was run through URET, a total of 443 conflicts were predicted to occur in one of the three selected sectors. Of those, 100 (23%) would not have been notified due to the application of probe notification logic. (The logic uses the probability of the conflict to determine the time to notify the controller.) For these 100 cases, the conflict was re-evaluated and determined to no longer exist prior to its notification time. The remaining 343 conflicts were split by class, 185 red and 158 yellow, then partitioned by status:

- 1) CLEAR: The controller issues a clearance to resolve the predicted conflict.
- 2) RECON: The trajectory is reconfirmed after notification of the conflict. The conflict is determined to no longer exist based on the new trajectory. No clearance is issued.
- 3) REMOD: The trajectory is remodeled after conflict notification. The conflict is determined to no longer exist based on the new trajectory. No clearance is issued¹.
- 4) EXPIRE: The conflict expires without any clearance issued.

Figure 2 shows the resulting disposition of conflicts by class and status.

These results are again consistent with the observations from the field trials. There is a high correlation between notified conflicts and controller actions (85% and 74% match for red and yellow problems, respectively). Red conflicts are more likely to result in a clearance while yellow conflicts are more likely to be closely monitored.

¹. A trajectory is remodeled upon entry of an amendment or some other specific event, e.g., entry of the first track report for the flight.




Figure 2. Disposition of Red and Yellow Conflicts from Voice Tape Analysis

Conflict Warning Time Some key operational objectives for the D-side conflict probe are to enable the sector team to smooth the sector workload, improve coordination, respond to user preferences, and improve the level of safety by providing more timely conflict information. To meet these objectives, adequate conflict warning time must be provided. Here, warning time is defined as the interval between the predicted conflict start time and the time at which the controller is notified of the conflict. URET functions determine a desired warning time that is between 10 and 20 minutes. The conflict probe notification

logic noted earlier calculates a warning time by interpolating between 10 and 20 minutes using the likelihood value associated with that conflict. Conflicts with a near-zero likelihood of conflict are set to notify close to a 10 minute warning time and conflicts with a likelihood of one are set to notify with a 20 minute warning time.

The results and analysis below use the warning times for unique conflicts that would have been notified to the controller (4,766 total conflicts). Since, in the scenario, some flights become available to the probe function with less than the desired warning time, e.g., departures or inbound aircraft from an adjacent facility, warning time data are categorized into two sets. The first set is comprised of conflicts initially predicted before the desired notification time. The second is for conflicts predicted after the desired notification time. The size of each set is 1,859 and 2,907, respectively.

For both red and yellow conflicts predicted before the desired notification time, the average warning time is almost 14 minutes. Conflicts predicted after the desired notification time have an average warning time of roughly 6.5 minutes. When these two sets are combined, the average warning time for all conflicts is approximately 9.5 minutes.

With the provision of more accurate and timely flight information in an interfacility operation², the number of yellow conflicts that are likely to turn red will tend to decrease. Based on this and the fact that controllers place a higher priority on red conflicts, warning time estimates for only those conflicts are provided. In the scenario, there are 792 unique red conflicts predicted before the desired notification time. Over 95% have a warning time of 10 minutes or more. Some have a warning time that is less than desired. These involve aircraft that are reconfirmed or amended between the first probe time and notification time such that a conflict originally predicted to occur disappears but then re-appears at a later time, when an immediate notification (of less than 10 minutes) is necessary. Figure 3 provides a distribution of warning times for these conflicts.

².Interfacility capabilities will allow an upstream URET system to provide better and earlier position and intent information to the downstream system.



Figure 3. Distribution of Warning Times for Red Conflicts

Again, considering the voice tape analysis, the time between alert notification and controller clearance was determined for each case. For both red and yellow conflicts, a total of 274 clearances were issued. The time in question was determined for 191 of those. The remaining 83 clearances did not come from any of the three sectors or they were delayed clearances. Figure 4 presents a distribution of the times for those 191 cases. The data show the mean time to be about 5 minutes.




Figure 4. Distribution: Time Before Clearance that Controller Would Have Been Notified of the Conflict

Achieving These Results A key finding from these analyses is that, by far, the most significant inaccuracies come from the uncertainties and constraints on the ATC environment in which URET operates. The effectiveness of the probe is constrained by data dependencies and limited knowledge of both controller and pilot intent. These can be grouped into four areas:

- 1) Facility Adaptation - undocumented or outdated procedural restrictions, data limited to Host flight plan processing requirements, inconsistencies between facilities, etc.
- 2) ATC Operations - dynamic changes to procedural restrictions, unknown control actions (e.g., vectors not entered as route amendments), actions with uncertain consequences (e.g., interim altitude clearances)
- 3) Aircraft Performance - unknown operator preferences, pilot and navigational variation, unknown equipment types and configurations, other unknowns (e.g., takeoff weight)
- 4) Predictive Uncertainty - wind error, missing or bad track data, lack of sophisticated tracking function, etc.

A number of techniques are used by URET in order to compensate:

- 1) Automated correction and expansion of Host adaptation.
- 2) Automated building and testing of procedural restrictions.
- 3) Adaptation of detailed aircraft performance characteristics.
- 4) Adaptive modeling of speeds and altitude gradients.
- 5) Speed and gradient estimation algorithms to alleviate surveillance and tracking inaccuracies
- 6) Heuristics to infer missing clearance amendments.
- 7) Heuristics to map unknown aircraft types to their closest match for data that are available.
- 8) Reasonableness checks on position and altitude updates.
- 9) Probabilistic conflict notification logic.

Overall, the approach of combining the lab analyses of real-world data with field trial results provided a useful mechanism to

validate the core URET functions and to refine them to the level of being operationally acceptable for problem solving and planning at the sector.

benefits analysis

Preliminary results from analyses of Conflict Probe benefits to controllers and pilots were presented in Reference 1. Here, we present results from more recent analysis.

Restrictions in Today's ATC System

In order to discuss the benefits to users from Conflict Probe, it is necessary to describe the restrictions to flight in today's ATC system. The FAA publishes a set of Preferred Instrument Flight Rules (IFR) High and Low Altitude Routes that must be flown by users flying between certain airports or flying through certain airspaces. These are referred to as ATC preferred routes. Many of these routes specify the entire route of flight, including departure and arrival routing. Most of the ATC preferred routes are associated with the major airports in the eastern U. S. There are very few for western airports. Some of these ATC preferred routes are rather close to direct routes between the associated airports, while others incorporate some degree of indirect routing in order to segregate major departure and arrival flows. However, the major effect of these routes is to force the users to fly the specified route every day, even though the wind-optimal route would vary from day to day and on occasion could differ substantially from the ATC preferred route.

The FAA also publishes Standard Terminal Arrival Routes (STARs) and Standard Instrument Departures (SIDs) that specify the route to be flown for arriving at or departing from major airports. These specify the routes to be flown within about 100 nm of the airports, and they apply to all aircraft, not just those subject to ATC preferred routes.

In addition to the preceding, FAA establishes a number of altitude and routing restrictions. Typically these specify that a controller must direct an aircraft arriving at a certain airport over a specified fix, or must clear that aircraft to a specified altitude when it crosses a particular sector or ARTCC boundary. These restrictions are established through Letters of Agreement between different ATC facilities and are reflected in the Standard Operating Procedures written for each sector. The altitude restrictions usually require an aircraft to descend to one or more lower altitudes than the desired cruise altitude and to spend some time in level flight at those lower altitudes.

FAA establishes one-way airways in congested traffic areas. These routes segregate major flows of traffic (often separating departure flows from arrival flows). These routes are prevalent in the airspace between Atlanta, Washington D.C., Cleveland, and Boston.

FAA applies altitude-for-direction rules in determining allowed cruising altitudes. For aircraft that are flying with a heading between 0 and 179 degrees, the cruising altitude must be FL 290, 330, 370, or 410. For those with heading of 180 to 359 degrees, the cruising altitude must be FL 310, 350, 390, or 430. Pilots flying an hour or more in the cruise phase of flight will want to go to higher altitudes as their aircraft burns fuel. With these altitude-for-direction rules in effect they must wait to climb to a higher altitude until they can climb 4,000 feet, if above FL290. The most efficient cruise profile for jet aircraft is a cruise climb. A power setting just slightly higher than that required for level flight is applied, and the aircraft is allowed to climb at a very low rate (typically less than 100 feet per minute). With 4,000 feet steps required, the step profile is a rather crude approximation to the desired cruise climb profile. Without the altitude-for-direction rules, 2,000 feet steps could be made, and a better approximation to cruise climb could be realized. If the requirement to cruise in level flight could be eliminated altogether the users could realize all of the potential benefit of cruise climb.

Many restrictions of the preceding types have been applied over the years as a means of providing structure to aid the controllers in manually handling growing levels of air traffic. Taken together, these restrictions have caused users to burn more fuel and to fly longer than if they could fly unrestricted. Airspace users in Reference 2 have expressed their very strong desire to reduce many of these restrictions as soon as possible.

National Route Program

FAA initiated the National Route Program (NRP) to allow aircraft to fly user preferred routes instead of the ATC preferred routes in certain circumstances. Since the initiation of this program, the airspaces where aircraft are permitted to fly NRP have been gradually expanded and the qualification rules somewhat relaxed. The program is described in Reference 3.

Currently, pilots can request NRP for flights that have a cruise altitude at or above FL290. However, the flight must be flight planned on the ATC preferred route within 200 nm of the departure airport and within 200 nm of the destination airport. Operators of regularly scheduled flights can request a standing pre-approval of NRP for certain departure/destination airport pairs. Requests for other NRP flights must be made on an individual flight basis. Approval of these routes will be made by coordinating with the affected ARTCCs for the specific flight. The request for the NRP route may be granted or denied. This gives the FAA flexibility to require that the flight be made on the ATC preferred route, if unusual circumstances (e.g., severe weather, ATC facility outage, navigation station outage) are forecast or being experienced

Most of the national airlines in the U.S. have provisions to utilize NRP for at least some of their flights, and they have reported substantial benefits from this program. However, because NRP is not granted for some of their flights, and because not all of their aircraft are equipped to take full advantage of NRP, a substantial number of flights still fly the ATC preferred routes.

Wind-optimal Route Analysis

The objective of this analysis was to determine the potential dollar savings that users in a real-world traffic scenario would experience if permitted to fly the wind-optimal horizontal route and preferred altitude profile instead of the route flown within the current airspace system. This is the potential benefit if every user were to fly wind-optimal routes and optimal profiles unrestricted. We do not claim and do not expect that use of a Conflict Probe will allow users to experience 100% of these potential benefits. Later in this paper we will discuss specific features of Conflict Probe, along with rationale, that we believe will enable users to realize some portion of the potential benefits.

This analysis also assessed the potential benefit if all aircraft that were qualified were allowed to fly NRP, and quantified the incremental benefit over that from NRP if all aircraft could fly unrestricted.

This analysis was carried out on a 24-hour U.S. nationwide traffic sample from 3 May 1995. The traffic sample was obtained from the FAA's Enhanced Traffic Management System (ETMS) that reports aircraft flight plan and track data from each en route facility to a central computer system once every five minutes. Forecast winds in effect at the time of this sample were captured and used for this analysis. General aviation and military flights and a few flights with erroneous or incomplete data were eliminated. About 21,000 flights remained in the analysis.

The analysis first took the data for the 21,000 flights and computed the time to complete the flight for each assuming that the flight made a continuous climb to the cruising altitude in the original flight plan and that the flight followed the horizontal flight path represented by the string of ETMS data points for that flight. The aircraft were assumed to cruise at the airspeed indicated in the original flight plan. No additional climbs were simulated. While the speed and altitude profiles for each flight did not match exactly those used in the actual flight, they served to establish a baseline flight time for the horizontal path actually flown. Using these assumed speed and altitude profiles for the baseline allowed us to use the same profiles for the wind-optimal route and the NRP route calculations. For these calculations the forecast winds were used to determine the ground speed for the flight in the baseline calculation.

A certain number of flights in the 3 May 1995 scenario were flying NRP routes. Their flight times in the baseline calculation were those calculated using the speed from the flight plan and the archived winds, not the times from the flights as flown.

The next step was to calculate new flight times for all aircraft that would have been qualified to fly NRP routes. Aircraft were qualified if the flight planned altitude was at or above FL290 and the distance from departure to destination airport was more than 750 nm. The points where the as-flown flight path in the baseline was 200 nm from the departure airport and the destination airport were established, and then the wind-optimal flight path between these two points was calculated. The total horizontal flight path for the NRP calculation was made up of the as-flown flight path on either end and the wind-optimal flight path in the middle. The wind optimization was used only to establish the horizontal flight path. Optimization in altitude was not performed. The aircraft flew the same altitude profile in this calculation as in the baseline calculation.

For the NRP calculation, it was assumed that ATC would permit all flights qualifying for NRP to fly the optimum route. This was done to provide an indication of the total potential benefit of NRP. It is understood that it is probably unrealistic to expect that all qualifying flights would be permitted to fly NRP routes in today's ATC system without new automation aids for the controller.

A third calculation was performed where all aircraft were permitted to fly the wind-optimal route from a point 20 nm from the departure airport to a point 20 nm from the destination airport. Again, the wind optimization was performed only to determine the horizontal flight path. The flight times were determined for this calculation. The altitude profile was the same as in the baseline calculation.

The time savings generated by the three calculations described above were converted to dollar savings by accounting for the cost of the fuel saved, and by including the savings in other direct operating costs resulting from the shortened flights. The results are presented in Table 1.

The 100% column indicates the dollar amount of the total potential benefits for the situation described in the first column. The first four rows present incremental benefits beyond those of the preceding row. The first row represents the benefits if all flights eligible for NRP are flying wind-optimal routes subject to the conditions of NRP. The last row shows the incremental benefits of totally unrestricted flight over those of NRP in the first row. The other columns simply show the dollar amount of savings that would result if the indicated percentage of the total potential benefits could be realized.

The results presented in the last four rows of Table 1 include benefits from some portion of each flight's path that occurs in terminal airspace, as well as a portion in en route airspace. A Conflict Probe provided to en route **Table 1. Annual Cost Savings in Millions of Dollars for all Airlines from NRP and Unrestricted Routes and Descents**

	Savings if a Percentage of Potential Savings is Realized			
	100 %	50 %	25 %	10 %
Additional Savings over Baseline if all Flights Eligible for NRP are on Wind-Optimal Horizontal Routes and Baseline Altitude Profiles	62			
Additional Fuel Savings Beyond the Preceding if all Flights Use Unrestricted Descents	105	53	26	11
Additional Fuel Savings Beyond the Preceding if all Flights use Wind-Optimal Horizontal Routes	255	128	64	26
Additional Other Direct Operating Cost Savings Beyond the Preceding	260	130	65	26
Total of Additional Cost Savings Beyond NRP (Sum of Preceding 3 Rows)	620	310	155	62

controllers is not expected to contribute to benefits in terminal airspace. A continuation of this analysis is partitioning the potential benefits into those achieved in terminal airspace and those in en route airspace; results of this partitioning are not available at this time.

We do not currently have a quantitative assessment of the percentage of the total potential benefits that the use of Conflict Probe could achieve. However, there was a qualitative assessment of this percentage provided by a team of controllers who interacted with AERA simulations and prototypes, and held numerous team meetings with AERA designers over a period of several years. The controllers estimated that 60% of the potential benefits could be realized through use of the full AERA capability (URET capabilities plus a conflict resolution capability). This was reported in Reference 4. To provide a conservative analysis, the authors of Reference 4 used a percentage of 40% in calculating the benefits of AERA.

We believe the percentage of potential benefits achievable from Conflict Probe is substantially greater than 10%, but even if it were as low as this Table 1 shows that the savings would be substantial and would amortize the expense of providing this capability in a very short time.

Later sections in this paper provide a description of the specific features of Conflict Probe that can help remove restrictions, and a qualitative discussion of why we believe these features can help achieve a meaningful percentage of the potential benefits. Additional study currently underway is attempting to provide a quantitative assessment of the percentage of potential benefits that might be achievable through interactive, human-in-the-loop simulations. In these simulations, subject controllers control a sector of traffic using the URET capabilities when a substantial number of aircraft are flown on wind-optimal routes.

ATC Preferred Route Analysis

The NRP results presented in the preceding section assumed that all flights eligible for NRP would fly wind-optimal routes. In this section we wish to determine how many aircraft are currently taking advantage of NRP, how many are still flying on ATC preferred routes, and what some of the ATC considerations are that might prevent greater use of NRP. Answering these questions provides an indication of the degree to which Conflict Probe can provide benefits beyond those of NRP.

A three-hour sample of real-world data from Indianapolis ARTCC was analyzed to answer these questions. The sample covered 1100Z to 1400Z (0600 to 0900 local time) on Wednesday, 11 December 1996. Traffic conditions were routine during this sample. All flight plans output by the Indianapolis Host computer during this sample were extracted. There were approximately 1,300 separate flights that passed through Indianapolis ARTCC airspace during this period. Within these flights there were 847 separate departure/destination airport pairs represented. This indicates how diverse the traffic in Indianapolis ARTCC is and indicates the degree to which the controllers must handle crossing traffic.

The flights were filtered to leave only those that had a flight planned altitude at or above FL290 and a stage length greater than 500 nm. There were 415 flights remaining after this filter. 329 of these flights were air carrier flights, 81 were general aviation, and 5 were military. Using the FAA published High Altitude Preferred Routes effective as of 5 December 1996, we determined whether or not each of these flights had a preferred route defined for its city pair, or whether it passed through airspaces where a preferred route was defined. 199 of these 415 flights had preferred routes defined for their routes.

We then analyzed each of these 199 flights individually to determine whether or not the flight was substantially following the prescribed preferred route. We did this by comparing the routing of the preferred route with the routing of the flight plan using a high altitude aeronautical chart. For this analysis we extracted and analyzed all of the flight plan and amendment messages for each aircraft that had been provided by the Host computer. In some cases the determination of whether the flight was on the preferred route or not was not clear - it could be declared either way. Generally, we considered the flight on the preferred route if a significant portion in the middle of the route was along the preferred route. We found that 127 of the 199 flights were on the preferred route, 66 flights were not, and 6 flights could not be determined, because we received only an abbreviated flight plan from the Host computer.

When the airline or the pilot files the flight plan for an NRP flight, the letters "NRP" are included in the remarks section of the flight plan. Of the 415 flights in our sample, 33 had NRP in the flight plan. It is possible that there were a few more aircraft than this participating in NRP. When a controller amends the remarks section of the flight plan, the NRP indication is removed, unless the controller includes it in his amendment.

All of those flights with NRP in the remarks section were air carrier flights. Each of the U.S. national passenger carriers had at least one flight in this sample that had NRP in the flight plan. This is a good indication that the airlines are well motivated to avoid ATC preferred routes and achieve the advantages of free flight. Since only a fraction of the air carrier flights are flying NRP now, the airlines are undoubtedly motivated to get more of their flights off the ATC preferred routes.

There was no evidence in this sample that general aviation aircraft participate in NRP. However, many of the general aviation flights take place between airports for which no preferred routes are specified, and there is ample evidence that the general aviation flights were filed on user preferred routes in these cases. There were also a number of general aviation flights for city pairs that did have ATC preferred routes defined, and for a number of these flights were not on the ATC preferred route. Evidently, some of the general aviation flights are being permitted to file and fly routes off the preferred routes, even without going through the NRP procedures.

There were individual flights in the sample that exhibited nearly every possible type of behavior with respect to ATC preferred routes. There was a small number of cases in which the original flight plan for a flight gave a route off the preferred route, and included NRP in the remarks, but a later amendment generated a flight plan that exactly followed the preferred route. There were a few cases where NRP was in the remarks, but the filed route was exactly the preferred route. There were a number of cases in which controllers issued direct routes to aircraft that were on the ATC preferred routes, presumably allowing the flight to fly according to user preference for at least a portion of the flight.

The overall picture that emerges from this analysis is the following. For about half the high altitude/long distance flights through Indianapolis ARTCC there are no ATC preferred routes defined. For those where there are preferred routes defined, about two-thirds of the flights are flying the preferred routes. While the users may be content to fly the preferred routes for some of these flights, there is evidence that the ATC system has required them to fly the preferred routes for a number of other flights. Two specific cases are discussed.

In Reference 1, a situation observed in an Indianapolis ARTCC traffic sample from 29 March 1995 was described. A cluster of eleven aircraft bound for Dallas/Ft. Worth, and all on airway J6 were observed in the vicinity of Charleston, WV. This situation is shown in Figure 5.

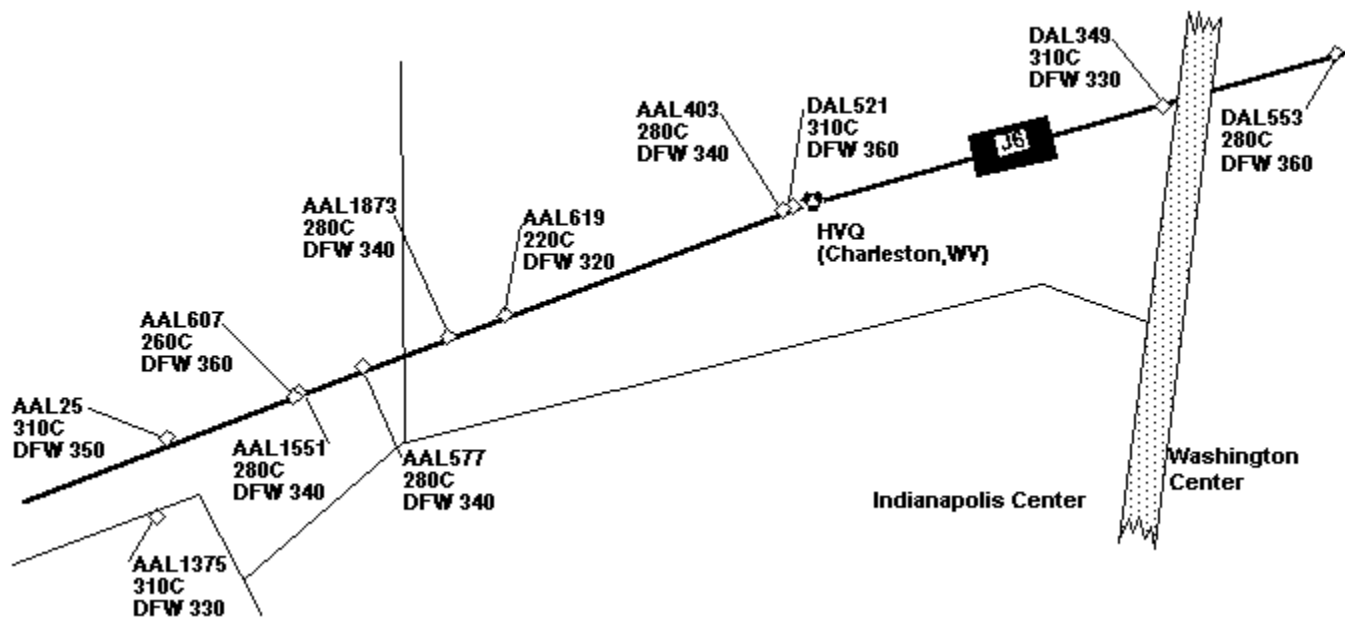


Figure 5. Eleven Aircraft Bound for DFW on the Same ATC Preferred Route

These eleven aircraft took off from eight different airports in the northeast - Boston, Bradley, Laganardia, Newark, Philadelphia, Baltimore/Washington, Washington National, and Dulles. It is surprising to find that these aircraft are essentially sequenced for arrival at DFW over Charleston, which is more than 800 miles from DFW. Each of the eight departure airports had a preferred route specified for DFW, the entire route was spelled out, and the route coincided with J6 through Indianapolis ARTCC. Using a plot of the forecast upper winds for the time of that sample, we constructed by hand a route for the flight from BOS to DFW that appeared to be more efficient than the preferred route. We submitted both the preferred route and the alternative to the URET system, and compared the times of arrival. Even though the alternative route had a ground track 29 miles longer, it had a flight time 4 minutes and 33 seconds shorter than the preferred route. This represents a 2% savings in flying time. Had we generated a true wind-optimum route, the time savings might have been greater. The Preferred IFR Route and the alternative are shown in Figure 6.

The qualification rules for NRP were made less restrictive after 29 March 1995, so we looked at other traffic samples from later dates to determine if NRP were used for these flights at a later time. Five different scenarios all from dates after 29 March 1995 were viewed, and every time there was a cluster of aircraft flying from various airports in the northeast to DFW, all the aircraft in the cluster were on the preferred route. To date we have not seen a single aircraft from any of these airports to DFW that was flying off the preferred route. The forecast winds for each of these scenarios had quite different patterns and it was apparent that the wind-optimal routes would have been quite different for each.

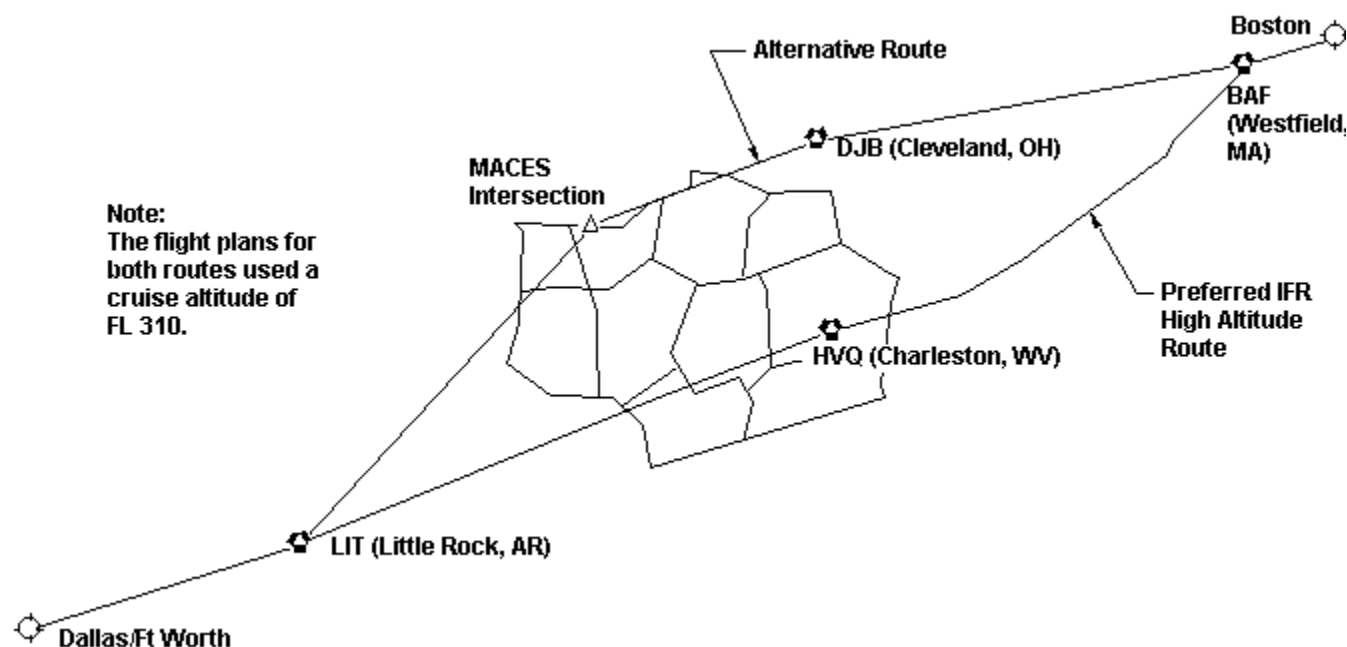


Figure 6. Preferred IFR High Altitude Route and Alternative for Flight from BOS to DFW

The second case was observed in the 11 December 1996 Indianapolis scenario. While analyzing the individual flights that had preferred routes defined, we observed that a great many of the preferred routes passed through the Rosewood VOR (identifier ROD). Figure 7 shows a snapshot of traffic in the vicinity of ROD at 1219Z. The positions and headings of 12 aircraft whose flight plans pass through ROD are shown as dots with arrows attached. The flights' call signs, current and assigned altitudes, and destination are given in data blocks for these aircraft. Circles with current and assigned altitudes without heading vectors are other flights of concern to the Dayton sector that do not pass over ROD. The sector boundaries of the Dayton sector are indicated. The departure and destination airports for each of the flights passing over ROD are shown as text. Those that are on an ATC preferred route are indicated. Included in this figure is a list of other departure/destination airports for which the ATC preferred route passes over ROD, but for which no flights were present in this snapshot.

It is logical and to some degree unavoidable that some of the flights whose routes originate or terminate in the vicinity of ROD should pass over ROD. But a number of the preferred routes involve departure/destination pairs where both airports are at some distance from ROD and these preferred routes might be able to be relaxed with use of Conflict Probe.

It is apparent that there is a great deal of traffic converging on ROD from a number of directions. Also, there is a merging of overflight traffic with traffic climbing from or descending to airports within about 150 nm of ROD. Our observation was that few of the aircraft whose ATC preferred routes passed through ROD were permitted to fly another route. The airspace around ROD is obviously a challenging ATC environment. Apparently, the ATC managers believe it is necessary to maintain the ATC preferred route structure in this area. With the use of Conflict Probe, some of this structure might be relaxed such that many of the flights could avoid passing over ROD. It is possible that this would actually reduce the net number of conflicts that controllers would have to resolve. Future work will investigate this possibility.

Restrictions in the Current ATC System and Rationale for Reducing or Removing them with Conflict Probe

Our involvement with the field evaluations and the analyses described here have allowed us to identify several types of restrictions in the ATC system and to understand how they help the controller handle traffic in the current environment. Each is discussed below. For each we describe the features of Conflict Probe that can support reducing or eliminating these

restrictions. Most of those features are included in Delivery 2.1 of URET that is now in place at Indianapolis and Memphis ARTCCs. Those features that are planned but

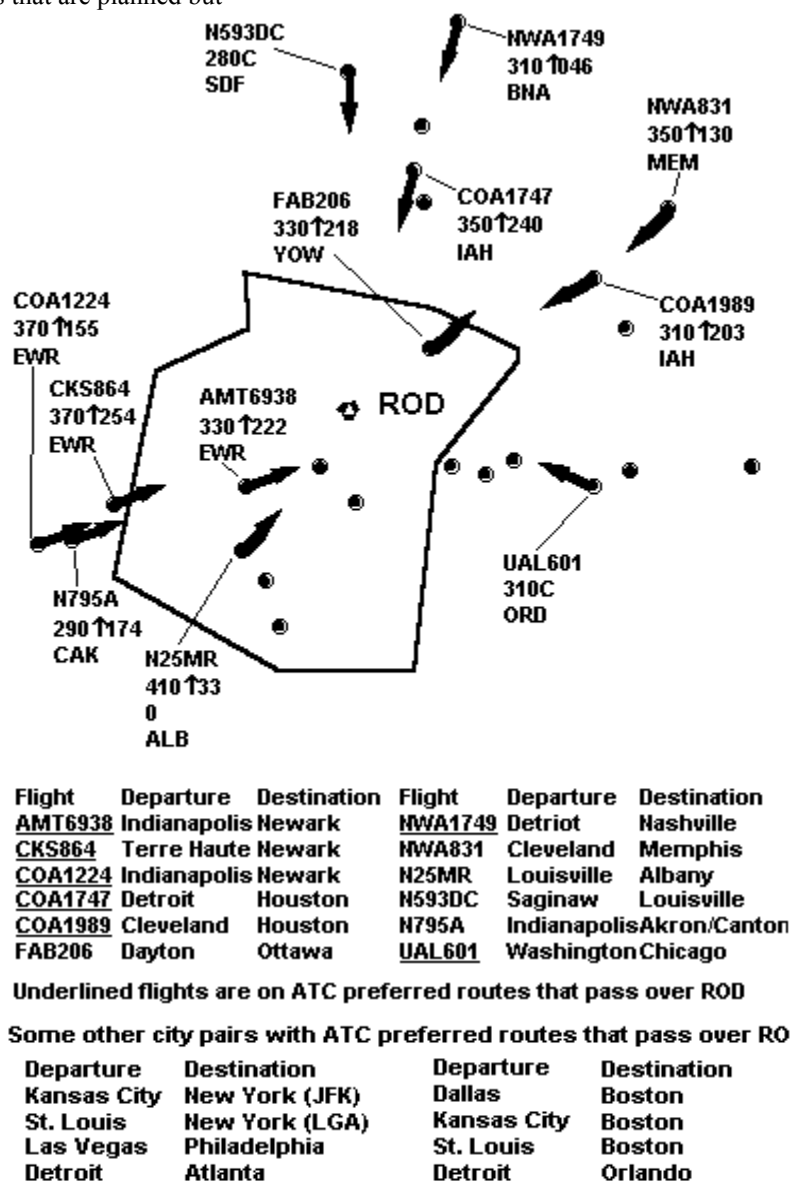


Figure 7. Snapshot of Flights Whose Flight Plans Pass Over ROD VOR

not yet included in Delivery 2.1 are so identified. Most of the Conflict Probe features would need to be implemented in a fault tolerant computer system in order to support removal of the restrictions. The URET prototype at these two ARTCCs does not have fault tolerance. However, the FAA has initiated activity toward the development of a production version of Conflict Probe that will implement the features of URET Delivery D2.1 in a fault-tolerant architecture.

We are not in any way being critical of the FAA for having these restrictions in place in the current ATC system. In our view, imposing these restrictions has been prudent and necessary to allow the controllers to handle the complex and heavy traffic situations with today's manual methods.

ATC Preferred Routes to Provide Predictability Requiring aircraft to fly ATC preferred routes provides predictability for the controllers and minimizes the amount of sector-to-sector coordination required. When most aircraft fly ATC preferred routes on airways, the ATC managers can arrange the sector boundaries with respect to these airways so that handoffs occur cleanly at sector boundaries. Figure 8 shows some of the high altitude sector boundaries and the jet airways in a portion of Indianapolis ARTCC. The degree to which the sector boundaries have been established in reference to the airways is quite apparent. One can see the ambiguity and coordination difficulty that can occur at sector boundaries when users file random

routes off airways from Figure 6. The alternative route lies right along a sector boundary for a considerable distance, and cuts just small corners of several sectors. Any conflicts encountered during these portions would require coordination with one or more other sectors. It is not impossible for controllers to handle aircraft on such flight plans; they do so often today. However, doing so adds an incremental workload that, when added to the other workload of heavy and complex traffic situations, makes control of traffic challenging.




Figure 8. Indianapolis ARTCC High Altitude Sectors and Airways

URET detects all conflicts without regard for the alignment of routes with respect to sector or ARTCC boundaries. The conflict notification feature of URET assigns responsibility for the resolution of each conflict to one and only one sector, so ambiguity about who is responsible for conflicts at a boundary are removed. All affected sectors are made aware of the conflict and they can determine which sector has been assigned responsibility. The autocoordination function of URET provides a convenient way for the responsible controller to notify another sector about the proposed resolution, and to request that sector to take action in those cases where the aircraft is currently under control of that other sector. The autocoordination function involves less workload than current voice coordination. Autocoordination is silent, non-interfering, and asynchronous, in contrast to voice coordination which requires both the initiator and the receiver to be on the interphone at the same time.

The current ATC automation system is built on the concept of fixes and fix posting areas. All aircraft whose flight plans pass through a given fix posting area will have a paper flight strip printed for the fix defined for that fix posting area. The fixes are usually VORs. There are usually one or two fixes per sector and the automation system will print the strips for a given fix at the sector that controls that fix. These posting fixes are easily identified on Figure 8 because they are the locations where a number of different airways intersect. Printing strips for the posting fixes serves as a form of strategic planning aid for the sector team. The strip has the current altitude and the current estimate of the time of arrival for the subject aircraft at that fix. By scanning the strips the team can identify potential future conflicts. When a significant number of aircraft are flying random routes and not passing over the posting fixes, the paper strips have less value as a strategic planning aid.

URET detects conflicts and displays the location of each conflict graphically, regardless of whether aircraft pass over posting fixes. Furthermore, URET's estimates of times along the trajectory are substantially more accurate than those represented on the paper flights strips. The Host computer only resynchronizes the times on the strips if the radar data indicates the arrival time will be off by more than 3 minutes. URET resynchronizes the trajectories if the times are off by about 15 seconds. Obviously, the URET trajectory modeling, conflict probe and graphical display capabilities provide a much more effective strategic planning capability than the current paper strips, and they do not require aircraft to pass over the posting fixes.

By requiring aircraft to fly on ATC preferred routes, the ATC system can control where many of the conflicts occur - they will occur at the posting fixes. To further make the task of resolving conflicts easier, and to reduce the amount of

coordination for the resolution of conflicts, the sector boundaries are arranged so that the posting fixes usually are in the interior of sectors and not close to the boundaries. This practice is also evident in Figure 8.

URET can detect conflicts occurring at any point and provides the notification logic and autocoordination logic that allow the controller to conveniently resolve conflicts occurring anywhere. URET does not require that aircraft fly ATC preferred routes so that most conflicts will occur cleanly in the interior of sectors.

Sector size is a factor that determines the amount and nature of the controller's workload. In the past when traffic growth led to difficult conditions in a particular sector, sector boundaries were redrawn in an effort to partition the workload. This usually led to smaller sectors. There is a limit to how small the sectors can be made. In smaller sectors, the controller spends a higher percentage of his time handing off aircraft and coordinating with other sectors. With smaller sectors, there is very little airspace in which to maneuver aircraft. Nearly every resolution requires coordination with another sector.

Figure 9 illustrates the point with a snapshot of a conflict that was detected by URET in the 11 December 1996 scenario. The conflict has been notified to Sector 87, the sector in which the predicted conflict will occur. However, neither aircraft is currently controlled by Sector 87, and in fact each is controlled by a different sector. At the time of the snapshot, the conflict will begin 14 minutes in the future. It is apparent that the size of the sector is smaller than the dimensions of many strategic conflicts that URET detects and displays to the controller. The small size of today's sectors encourages controllers to resolve conflicts later and more tactically, because the additional coordination necessary to resolve them earlier would increase their workload.




Figure 9. An Alert Notified to Sector 87 14 Minutes Before the Start of Conflict

The URET strategic conflict probe, notification, trial planning, display, and autocoordination features allow controllers to deal with conflicts strategically even with today's small sector sizes.

The predictability of having aircraft on ATC preferred routes helps ATC managers with their planning in a number of ways. Peaking patterns will repeat from day to day. This allows the area supervisors to plan breaks, shift changes, and the decombining of sectors in anticipation of the peaks. It is obviously easier and safer to carry out shift changes and resectorization in advance of such peaks than in the midst of an unanticipated workload peak. The supervisor learns which sectors are more easily handled and which are more challenging at particular times, and can take this into account when assigning less experienced controllers to sectors. When a significant number of aircraft are flying wind-optimal routes, the routes will change from day to day, and some amount of this predictability will be lost.

The URET capabilities can help controllers deal safely with more challenging situations, so the controllers are less vulnerable to unanticipated peaks. Requiring aircraft to fly on ATC preferred routes in order to preserve the current degree of predictability would not be required. Supervisors can monitor the controller's URET display and get an instant indication of the number of strategic conflicts for which that sector is responsible. Because these conflicts are notified with alert times

substantially longer than are required to resolve conflicts tactically, supervisors have an opportunity to recognize a developing peak situation with time to take some remedial action.

A future capability that can be incorporated into a Conflict Probe is a workload alert. URET models trajectories for most aircraft all the way to the destination airport. It is possible to extend the lookahead time for the conflict probe function beyond the times at which alerts are notified to controllers. So such a workload alert function could detect potential conflicts 20 to 30 minutes in the future. An alert mechanism could be programmed to deliver an alert to the supervisor whenever the number of predicted conflicts exceeds a designated threshold. Alerts could also be generated for other workload measures.

Today, requests to fly NRP routes instead of ATC preferred routes for many of the routes are made by contacting the ATC System Command Center (ATCSCC) for each individual flight. Staff at the ATCSCC and at the ARTCCs affected by the proposed route decide to approve or deny the route. Because this decision is made perhaps 3 or 4 hours before the aircraft will fly through the airspace, there must be a prediction of the conditions that will exist in the future. Since it is difficult to predict traffic or weather conditions very precisely that far in advance, the FAA personnel may deny a request simply as a defensive measure. If approvals were granted routinely to fly NRP routes instead of certain ATC preferred routes, then it would seem that those particular routes could be eliminated as ATC preferred routes. Evidently, having the ability to deny NRP flights on some occasions while approving them on others is important to the FAA managers and controllers.

The features of URET listed above that allow controllers to more readily handle unanticipated situations should make it possible for FAA personnel to approve NRP routes more often, and possibly even to eliminate a number of the ATC preferred routes.

SIDs, STARs, One-Way Routes, Arrival and Departure Restrictions to Segregate Flows of Traffic These types of restrictions are used to segregate major flows of traffic. They utilize floors and ceilings and arrival and departure corridors to thread flows coming from different quadrants to various runway ends. The approach is to prevent the intersection of flows to the extent possible. This allows the controllers to concentrate on merging traffic in a given flow and to resolving conflicts between flights within a flow, rather than dealing with conflicts between aircraft in different flows.

Although these restrictions are created to organize flows to and from specific airports, their effect can extend a considerable distance from the airport. For example, the Little Rock transition to the Bonham 2 arrival to Dallas Fort Worth, begins at Little Rock which is about 250 nm from the airport. We do not expect that the use of Conflict Probe would eliminate these types of restrictions entirely. However, it is thought that their use could be confined to areas much closer to the airports, so that users could fly user preferred trajectories over a greater portion of their flight.

Many of the routes and restrictions are used to reduce the need for sector-to-sector coordination as much as possible. Today, if the controller contemplates changing the clearance to an aircraft, either to resolve a conflict, or to respond to a pilot request, he will as a practice coordinate with a downstream sector. This is especially true if the change would deviate from a SID or STAR or a restriction. In the future, it is possible that through the use of URET this practice could be reduced. The procedure might be changed for many situations to one that requires the controller to perform a trial plan on the proposed clearance, and allows him to issue the clearance without coordinating if the trial plan is conflict free. Even if it is deemed desirable to coordinate the proposal, this can be done more readily using the autocoordination function than with voice coordination. By eliminating the need for coordination, or by making it easier to coordinate, the use of Conflict Probe should help eliminate some routes and restrictions, and confine the extent of others to areas closer to the airports.

Altitude-for-Direction Rules to Provide More Scan Time

The altitude-for-direction rules were put in place to provide a safety margin for controllers who must visually scan the radar screen searching for future conflicts. Without these rules, controllers could often have head-on conflicts with aircraft at the same altitude. If the controller were to become momentarily distracted with another task, there might be very little time left for resolution once he recognizes the conflict. Imposing these rules means that aircraft cruising at the same altitude will very rarely have head-on conflicts, and the conflicts they do have will have modest closing speeds and allow more time for the controller's scan to detect the conflict.

With a dependable Conflict Probe in operation there is no need for the altitude-for-direction rules. Conflicts will be declared with nearly the same warning time regardless of closing angle. Furthermore, timely conflict recognition is more certain with the visual alerting of Conflict Probe than with the scanning detection of the controller.

If the altitude-for-direction rules were eliminated to permit opposite direction same altitude traffic, it would not be a big step to further relax the cruise altitude rules and permit cruise climbs. It would not be substantially more difficult for Conflict Probe to detect conflicts involving aircraft in cruise climb than it would be for detecting alerts involving level aircraft. URET currently does not have a feature in the trajectory modeling software to model cruise climbs, but such a capability could be easily added.

references

- [1] D. J. Brudnicki, A. L. McFarland, S. M. Schultheis, 1996, *Conflict Probe Benefits to Controllers and Users - Indications from Field Evaluations*, MP96W194, The MITRE Corporation, McLean, VA.
- [2] RTCA Task Force 3, 1995, *Final Report of RTCA Task Force 3 - Free Flight Implementation*, RTCA, Inc., Washington, D.C.
- [3] Federal Aviation Administration, 1996, *National Route Program*, Advisory Circular Number 90-91A, Federal Aviation Administration, Washington, D.C.
- [4] Butcher, A. M., Cheslow, M. D., Nelson, G. G., Reid, D. H., Russell, H. E., Sinnott, J. H., Walker, G.R., 1987, *The Advanced Automation System: A Benefit/Cost and Risk Analysis, Vol. IV, Appendices*, MTR87W235-04, The MITRE Corporation, McLean, VA.